

US009158591B2

(12) **United States Patent**
Matousek

(10) **Patent No.:** **US 9,158,591 B2**
(45) **Date of Patent:** **Oct. 13, 2015**

(54) **SYSTEM AND METHOD FOR CONTROLLED SHARING OF CONSUMABLE RESOURCES IN A COMPUTER CLUSTER**

(71) Applicant: **Dell Products L.P.**, Round Rock, TX (US)

(72) Inventor: **Jon Franklin Matousek**, Austin, TX (US)

(73) Assignee: **Metric Holdings, LLC**, Austin, TX (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 394 days.

(21) Appl. No.: **13/690,495**

(22) Filed: **Nov. 30, 2012**

(65) **Prior Publication Data**

US 2014/0115595 A1 Apr. 24, 2014

Related U.S. Application Data

(60) Provisional application No. 61/717,798, filed on Oct. 24, 2012.

(51) **Int. Cl.**
G06F 9/46 (2006.01)
G06F 9/50 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 9/5061** (2013.01); **G06F 9/5011** (2013.01); **G06F 2209/5014** (2013.01)

(58) **Field of Classification Search**
USPC 718/103
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,219,347 B1 5/2007 Waddington
7,725,583 B2 5/2010 Jackson

8,286,179 B2 10/2012 Comer
8,612,990 B1 12/2013 Greenfield et al.
2003/0028656 A1* 2/2003 Babka 709/229
2007/0094665 A1* 4/2007 Jackson 718/104
2007/0256077 A1* 11/2007 Zhong 718/104
2009/0012930 A1* 1/2009 Jackson 707/2
2010/0262964 A1 10/2010 Uyeda et al.
2011/0072138 A1 3/2011 Canturk et al.
2011/0314477 A1 12/2011 Zhong
2012/0079495 A1* 3/2012 Conti et al. 718/104
2012/0079501 A1 3/2012 Sandstrom
2012/0198462 A1* 8/2012 Cham et al. 718/103
2012/0221886 A1 8/2012 Barsness et al.
2013/0047162 A1 2/2013 Stefanov et al.

FOREIGN PATENT DOCUMENTS

EP 1318453 A1 6/2003

OTHER PUBLICATIONS

Sin, Chandandeep, "Completely Fair Scheduler", Linux Journal, Aug. 1, 2009, 7 pages.

(Continued)


Primary Examiner — Emerson Puente

Assistant Examiner — Steven Do

(57) **ABSTRACT**

In one embodiment, a method includes empirically analyzing a set of active reservations and a current set of consumable resources belonging to a class of consumable resources. Each active reservation is of a managed task type and includes a group of one or more tasks requiring access to a consumable resource of the class. The method further includes, based on the empirically analyzing, clocking the set of active reservations each clocking cycle. In addition, the method includes, responsive to the clocking, sorting a priority queue of the set of active reservations.

20 Claims, 16 Drawing Sheets



(56)

References Cited

OTHER PUBLICATIONS

<http://doc.opensuse.org/documentation/html/openSUSE/opensuse-tuning/cha.tuning.tasks>, “Chapter 14. Tuning the Task Scheduler”, Nov. 30, 2012, 13 pages.

http://en.wikipedia.org/wiki/Completely_Fair_Scheduler, “Completely Fair Scheduler”, Aug. 30, 2012, 4 pages.

Li, Tong et al., “Efficient and Scalable Multiprocessor Fair Scheduling Using Distributed Weighted Round-Robin”, Feb. 14-18, 2009, 10 pages.

Jones, M. Tim, “Inside the Linux 2.6 Completely Fair Scheduler”, Dec. 15, 2009, 6 pages.

Dong, Fangpeng et al., “Scheduling Algorithms for Grid Computing: State of the Art and Open Problems”, Technical Report No. 2006-504, Jan. 2006, 55 pages.

Zaharia, Matei, et al., “Job Scheduling for Multi-User MapReduce Clusters”, Technical Report No. UCB/EECS-2009-55, Apr. 30, 2009, 18 pages.

Isard, Michael et al., “Quincy: Fair Scheduling for Distributed Computing Clusters”, Microsoft Research, 2009, 20 pages.

* cited by examiner




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7





FIG. 8






FIG. 10




FIG. 11




FIG. 12




FIG. 13


1400
↘

FIG. 14




FIG. 15




FIG. 16

1

SYSTEM AND METHOD FOR CONTROLLED SHARING OF CONSUMABLE RESOURCES IN A COMPUTER CLUSTER

CROSS-REFERENCE TO RELATED APPLICATIONS

This patent application claims priority from, and incorporates by reference the entire disclosure of, U.S. Provisional Patent Application No. 61/717,798 filed on Oct. 24, 2012. This patent application incorporates by reference the entire disclosure of a U.S. patent application Ser. No. 13/690,526, filed on Nov. 30, 2012.

BACKGROUND

1. Technical Field

The present invention relates generally to scheduling and more particularly, but not by way of limitation, to systems and methods for controlled sharing of consumable resources in a computer cluster.

2. History of Related Art

A multi-tenant computer cluster such as, for example, a Software as a Service (SaaS) cloud, often uses computing resources (e.g., hardware and software resources) to perform services for customers. The computing resources can include, for example, computers (e.g., physical computers or virtual machines (VMs)), raw (block) and file-based storage, firewalls, load balancers, internet protocol (IP) addresses, virtual local area networks (VLANs), software bundles, and computing platforms that typically include an operating system, a programming-language execution environment, a database, and a web server. Services generally involve use of at least a portion of the computing resources for the benefit of the customer. The services can include, for example, emergency notification, accounting, collaboration, customer relationship management (CRM), management information systems (MIS), enterprise resource planning (ERP), invoicing, human resource management (HRM), content management (CM), service-desk management, and/or many other services.

Oftentimes, the multi-tenant computer cluster receives customer requests for service randomly responsive to needs that arise unpredictably. In addition, the customer requests for service frequently spawn other requests for service. Therefore, all requests are traditionally managed in a queue and serviced in a first-in first-out (FIFO) manner. As the queue becomes backlogged, the FIFO manner of servicing results in an unequal distribution of the computing resources across customers. Because the requests are serviced sequentially, customers with more requests are awarded a greater share of the computing resources than customers with fewer requests. The distribution of the computing resources across all customers is not generally controllable.

Moreover, as the value and use of information continues to increase, individuals and businesses seek additional ways to process and store information. One option available to users is information handling systems. An information handling system generally processes, compiles, stores, and/or communicates information or data for business, personal, or other purposes thereby allowing users to take advantage of the value of the information. Because technology and information handling needs and requirements vary between different users or applications, information handling systems may also vary regarding what information is handled, how the information is handled, how much information is processed, stored, or communicated, and how quickly and efficiently the information may be processed, stored, or communicated. The

2

variations in information handling systems allow for information handling systems to be general or configured for a specific user or specific use such as financial transaction processing, airline reservations, enterprise data storage, or global communications. In addition, information handling systems may include a variety of hardware and software components that may be configured to process, store, and communicate information and may include one or more computer systems, data storage systems, and networking systems.

SUMMARY OF THE INVENTION

In one embodiment, a method includes empirically analyzing, by at least one process on a computer cluster comprising a plurality of computers, a set of active reservations and a current set of consumable resources belonging to a class of consumable resources. Each active reservation is of a managed task type and comprises a group of one or more tasks requiring access to a consumable resource of the class. The method further includes, based on the empirically analyzing, clocking the set of active reservations each clocking cycle. In addition, the method includes, responsive to the clocking, sorting, by the at least one process, a priority queue of the set of active reservations.

In one embodiment, an information handling system includes a computer cluster comprising a plurality of computers. The computer cluster is operable to empirically analyze a set of active reservations and a current set of consumable resources belonging to a class of consumable resources. Each active reservation is of a managed task type and includes a group of one or more tasks requiring access to a consumable resource of the class. The computer cluster is further operable, based on the empirical analysis, to clock the set of active reservations each clocking cycle. In addition, the computer cluster is operable, responsive to the clocking, to sort a priority queue of the set of active reservations.

In one embodiment, a computer-program product includes a computer-usable medium having computer-readable program code embodied therein, the computer-readable program code adapted to be executed to implement a method. The method includes empirically analyzing a set of active reservations and a current set of consumable resources belonging to a class of consumable resources. Each active reservation is of a managed task type and includes a group of one or more tasks requiring access to a consumable resource of the class. The method further includes, based on the empirically analyzing, clocking the set of active reservations each clocking cycle. In addition, the method includes, responsive to the clocking, sorting a priority queue of the set of active reservations.

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of the method and apparatus of the present invention may be obtained by reference to the following Detailed Description when taken in conjunction with the accompanying Drawings wherein:

FIG. 1 illustrates a computer cluster;

FIG. 2 illustrates a flow-control scheme;

FIG. 3 illustrates a flow-control instance;

FIG. 4 illustrates a process that may be executed by a flow-control instance;

FIG. 5 illustrates a derivation of reservations from tasks;

FIG. 6 illustrates an exemplary balanced-utilization partitioning scheme;

FIG. 7 illustrates an implementation of a balanced-utilization partitioning scheme;

FIG. 8 illustrates an exemplary balanced-utilization partitioning scheme;

FIG. 9 illustrates an implementation of a balanced-utilization partitioning scheme;

FIG. 10 illustrates a balanced-utilization partitioning scheme;

FIG. 11 illustrates an implementation of a balanced-utilization partitioning scheme;

FIG. 12 illustrates a process for calculating an inner stationary distribution;

FIG. 13 illustrates clocking functionality of a flow-control instance;

FIG. 14 illustrates a process for decrementing a wait-time;

FIG. 15 illustrates a collection of interactive processes that may be executed by a flow-control instance; and

FIG. 16 illustrates a collection of interactive processes that may be executed by a flow-control instance.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS OF THE INVENTION

In various embodiments, customers can be served more equitably and controllably using systems and methods described herein. In a typical embodiment, flow control is instituted on a computer cluster by defining a class of consumable resources and establishing a framework governing utilization of the consumable resources by customers. For purposes of this patent application, a computer cluster is a set of loosely-connected computers, virtual or physical, that work together such that the computers can be viewed as a single system. In various embodiments, a computer cluster can provide a cloud environment such as, for example, a SaaS cloud environment.

A consumable resource, as used herein, refers to a limited resource that is accessible to a computer cluster. For example, consumable resources can include, but are not limited to, available memory, central processing units (CPUs), free space on a file system, network bandwidth, floating software licenses, voice-port hardware (e.g., text-to-speech voice ports) and access to a computing process. For simplicity, consumable resources may be referenced periodically herein as resources. In particular, with respect to access to a computing process, the computing process may be, for example, a bottleneck point in an overall business process. It should be appreciated that consumable resources can be either internal or external to a given computer cluster. It should further be appreciated that consumable resources can also be used to abstract human workflow. For example, in some embodiments, principles described herein are implemented in a call center that receives incoming calls to be serviced by call-center employees. In these embodiments, either the call-center employees or voice-communication channels staffed by such employees can be considered consumable resources that are controlled and managed as described herein.

A service, as used herein, is a semantic process or a combination of semantic processes that can be performed for the benefit of a customer. Services are generally requested by customers. A semantic process, as used herein, refers to one or more tasks performed by a computer. Tasks of a given task type, which can span multiple semantic processes, may utilize one or more consumable resources from a specific class of consumable resources. A class of consumable resources, as used herein, refers to an aggregation of like consumable resources that can fulfill a resource need, for example, of the given task type. In such cases, a flow-control instance can be utilized to control sharing of the specific class of consumable resources relative to tasks of the given task type. For example,

a class of consumable resource could include a batch of tasks or task groups that can be executed at a given time or allowed to access a computing process. A flow-control instance, as used herein, is software that is configured to manage and control sharing of a particular class of consumable resources among a plurality of customers.

FIG. 1 illustrates a computer cluster 100 that is operable to provide one or more services to customers. The computer cluster 100 includes a computer 102(1), a computer 102(2), a computer 102(3), and a database 106. For convenient reference, the computer 102(1), the computer 102(2), and the computer 102(3) may be referenced collectively as computers 102. In various embodiments, the computers 102 can be virtual computers, physical computers, or a combination thereof. For illustrative purposes, the computers 102 are shown to include the computer 102(1), the computer 102(2), and the computer 102(3). However, one of ordinary skill in the art will appreciate that, in practice, the computers 102 can include any number of physical and/or virtual computers. As shown, the computers 102 and the database 106 are operable to communicate over a network 108.

In a typical embodiment, each of the computers 102 is operable to execute one or more semantic processes related to the provision of the one or more services by the computer cluster 100. In particular, the computer 102(1) executes a process 104a(1), the computer 102(2) executes a process 104b(1) and a process 104c(1), and the computer 102(3) executes a process 104a(2), a process 104b(2), and a process 104c(2). In a typical embodiment, the process 104a(1) and the process 104a(2) are identical processes that have been replicated on each of the computer 102(1) and the computer 102(3), respectively. Likewise, the process 104(b)(1) and the process 104(b)(2) are identical processes that have been replicated on each of the computer 102(2) and the computer 102(3), respectively. Similarly, the process 104(c)(1) and the process 104(c)(2) are identical processes that have been replicated on each of the computer 102(2) and the computer 102(3), respectively.

Operationally, the one or more services provided by the computer cluster 100 can be initiated in a variety of ways. For example, in a typical embodiment, services can be requested by one of the customers via, for example, an interface provided over a public network such as, for example, the Internet. Once a service is initiated, the initiated service may encompass semantic processes selected from the processes 102(a)(1), 102(a)(2), 102(b)(1), 102(b)(2), 102(c)(1), and 102(c)(2). The semantic processes of the initiated service generally include tasks to be executed by the computers 102. In various embodiments, each service, semantic process, and task can spawn other services, semantic processes, and tasks, respectively, so that the initiated service results, for example, in many thousands of tasks. In a typical embodiment, some of those tasks may be of a task type that requires access to a specific class of consumable resources. In various embodiments, the computer cluster 100 controls sharing of the specific class of consumable resources via a flow-control instance that is executed on one or more of the computers 102. Examples of flow-control instances that may be utilized will be described with respect to the ensuing Figures.

FIG. 2 illustrates a flow-control scheme 200 involving process threads 216, 218, and 220. The flow-control scheme 200 illustrates that flow-control instances generally serve a gating function. The process threads 216, 218, and 220 each represent business logic that may be executed in parallel by a computer cluster such as, for example, the computer cluster 100 of FIG. 1, in the provision of one or more services to a customer. The process threads 216, 218, and 220 can each

5

include one or more semantic processes, which semantic processes can each further include one or more tasks.

The process thread **216** includes a logical path **216(1)**, a flow-control instance **216(2)**, and a semantic process **216(3)**. As shown, flow of the process thread **216** is gated by the flow-control instance **216(2)**. The flow-control instance **216(2)** controls access to a class of consumable resources needed by a task of the process thread **216**. Therefore, forward progress in the process thread **216**, and execution of the semantic process **216(3)**, is conditioned upon the flow-control instance **216(2)** providing consumable-resource access to the task and the task executing. After the task has executed, the semantic process **216(3)** can be initiated.

The process thread **218** includes a logical path **218(1)**, a flow-control instance **218(2)**, a flow-control instance **218(3)**, a flow-control instance **218(4)**, and a semantic process **218(5)**. As shown, flow of the process thread **218** is gated by the flow-control instances **218(2)**, **218(3)**, and **218(4)**. The flow-control instances **218(2)**, **218(3)**, and **218(4)** each control access to a distinct class of consumable resources needed by three distinct tasks of the process thread **218**. For purposes of illustration, the flow-control instances **218(2)**, **218(3)**, and **218(4)** may be considered to represent a first task, a second task, and a third task, respectively. Initially, forward progress in the process thread **218** is conditioned upon: (1) the first task being granted resource access by the flow-control instance **218(2)** and executing; and (2) the second task being granted resource access by the flow-control instance **218(3)** and executing. Therefore, as shown, the third task cannot execute until the first task and the second task are granted consumable-resource access by the flow-control instances **218(2)** and **218(3)**, respectively, and execute.

Once the first task and the second task have executed, forward progress in the process thread **218**, and execution of the semantic process **218(5)**, is conditioned upon the flow-control instance **218(4)** providing consumable-resource access to the third task and the third task executing. After the third task has executed, the semantic process **218(5)** can be initiated. In that way, the process thread **218** utilizes three flow-control instances in the provision of the one or more services.

The process thread **220** includes a logical path **220(1)**, a flow-control instance **220(2)**, and a semantic process **220(3)**. As shown, flow of the process thread **220** is gated by the flow-control instance **220(2)**. The flow-control instance **220(2)** controls access to a class of consumable resources needed by a task of the process thread **220**. Therefore, forward progress in the process thread **220**, and execution of the semantic process **220(3)**, is conditioned upon the flow-control instance **220(2)** providing consumable-resource access to the task and the task executing. After the task has executed, the semantic process **220(3)** can be initiated.

FIG. 3 illustrates a flow-control instance **300** that may be implemented on a computer cluster such as, for example, the computer cluster **100** of FIG. 1. The flow-control instance **300** performs controlled sharing **316** of a current set of consumable resources **310** for a given task type being managed by the flow-control instance **300** (hereinafter, "managed task type"). The current set of consumable resources **310** corresponds to a class of consumable resources required by the managed task type. The flow-control instance **300** enforces the controlled sharing **316** among a current set of active customers **312** according to a balanced-utilization partitioning scheme **314**. In a typical embodiment, the flow-control instance **300** is resident and executing on at least one computer of the computer cluster.

6

In a typical embodiment, the flow-control instance **300** manages tasks of the managed task type by grouping tasks into reservations. For example, each reservation can include those tasks that have a same task key. The task key usually defines a set of characteristics that justify grouping tasks into a same reservation. The set of characteristics can include, for example, a service ID for a requested service with which the task is associated, a customer ID corresponding to a requesting customer, a priority level assigned to the task, and the like. Reservations will be described in greater detail with respect to the ensuing Figures.

In a typical embodiment, the current set of consumable resources **310** includes a defined quantity of consumable resources for which the flow-control instance **300** is responsible for enforcing the controlled sharing **316**. The defined quantity can dynamically change during execution of the flow-control instance **300**. For example, consumable resources may become available or unavailable to the flow-control instance **300** and, correspondingly, be included in or excluded from the current set of consumable resources **310**. For illustrative purposes, the current set of consumable resources **310** is shown to include consumable resources **310(1)**, **310(2)**, and **310(3)**, which resources may be considered to represent an exemplary snapshot-in-time of the current set of consumable resources **310**.

The current set of active customers **312** includes customers that, at a given point in time, have at least one reservation of the managed task type that has not been processed to completion by the computer cluster (i.e., customers having active reservations). In a typical embodiment, the current set of active customers **312** changes dynamically during execution of the flow-control instance **300**. As new requests for service are received and corresponding reservations are created, customers who do not already have at least one active reservation of the managed task type are added to the current set of active customers **312**. As reservations are processed to completion, customers having no active reservations of the managed task type are removed from the current set of active customers **312**.

The balanced-utilization partitioning scheme **314** establishes how the current set of consumable resources **310** should be distributed to and consumed by the current set of active customers **312**. More particularly, the balanced-utilization partitioning scheme **314** defines logical partitions relative to a totality of the current set of consumable resources **310**. Each logical partition of the balanced-utilization partitioning scheme **314** is assigned a resource percentage representing a proportion of a given set of consumable resources (e.g., the current set of consumable resources **310** or a subset thereof) that, ideally, should be allocated to the logical partition. For example, in various embodiments, a total number of partitions may equal a number of active customers in the current set of active customers **312**. In various other embodiments, customers may be allowed to prioritize their tasks. In these embodiments, the total number of partitions could vary between the total number of active customers and the total number of active customers multiplied by a total number of priority levels enabled by the flow-control instance.

For illustrative purposes, the balanced-utilization partitioning scheme **314** is shown to include partitioning schemes **314(1)**, **314(2)**, and **314(3)**, which schemes may be considered an exemplary snapshot-in-time of the balanced-utilization partitioning scheme **314**. Like the current set of consumable resources **310** and the current set of active customers **312**, the balanced-utilization partitioning scheme **314** dynamically changes during execution of the flow-control instance **300** based on, for example, the current set of active customers **312** and a current set of active reservations of the

managed task type. As the balanced-utilization partitioning scheme **314** changes, new resource percentages are computed and assigned.

FIG. **4** illustrates a process **400** that may be executed, for example, by a flow-control instance that is resident and executing on at least one computer of a computer cluster. Various portions of the process **400** may be executed as part of the controlled sharing **316** of FIG. **3**. It should be appreciated that steps **402-410** of the process **400** are shown to execute sequentially for purposes of illustrating how tasks logically progress from being grouped into reservations to being assigned free resources. Although the process **400** can execute sequentially as shown, in a typical embodiment, steps **402-410** each represent subprocesses that execute in parallel on the computer cluster. In addition, as described below, various steps of the process **400** execute continuously to control sharing of a current set of consumable resources among a current set of active reservations. The process **400** begins at step **402**.

At step **402**, the flow-control instance derives reservations. Derivation of reservations involves grouping tasks that have, for example, a same task key, into reservations. As indicated, step **402** executes repeatedly due to constant creation of new tasks of a managed task type. Derivation of reservations will be described in greater detail with respect to FIG. **5**. After each execution of step **402**, the process **400** proceeds to step **404**. At step **404**, the flow-control instance empirically analyzes the current set of active reservations and the current set of consumable resources. In a typical embodiment, step **404** yields a balanced-utilization partitioning scheme such as, for example, the balanced-utilization partitioning scheme **314** of FIG. **3**. Examples of empirical analysis that can occur at step **402** will be described in greater detail with respect to FIGS. **6-12**. After execution of step **404**, the process **400** proceeds to step **406**.

At step **406**, the flow-control instance performs clocking of the current set of active reservations. In a typical embodiment, the flow-control instance maintains a virtual clock that, every clocking cycle, initiates a virtual pendulum for each reservation of the current set of active reservations. The clocking cycle can be, for example, a configurable unit of wall time (e.g., seconds or milliseconds) or a configurable number of CPU cycles. The initiation of the virtual pendulums results in each reservation being "clocked" according to a flow-control clocking weight that is specific to the reservation. Clocking of reservations will be described in more detail with respect to FIGS. **13-14**. As indicated, step **406** executes repeatedly for each clocking cycle. After step **406**, the process **400** proceeds to step **408**.

At step **408**, the flow-control instance determines whether a consumable resource of the current set of consumable resources is free (i.e., available for use). Whenever a consumable resource is free, the process **400** proceeds to step **410**. As indicated in FIG. **4**, once initiated, step **408** executes continuously as long as reservations exist in the priority queue and there are free resources available to assign. At step **410**, the flow-control instance assigns the free consumable resource to a most-needy resource in the priority queue. Examples of activities that can occur at steps **408** and **410** will be described in greater detail with respect to FIGS. **15-16**. After step **410**, the process **400** ends.

FIG. **5** illustrates a derivation **500** that can be used to derive reservations from tasks. The derivation **500** illustrates functionality that, in a typical embodiment, can be performed as part of step **402** of FIG. **4**. One or more flow-control instances **518** are shown to derive reservations from a task table **502**. For purposes of illustration, the one or more flow-control

instances **518** are shown to include a flow-control instance **518a**, a flow-control instance **518b**, a flow-control instance **518c**, and a flow-control instance **518d**. In a typical embodiment, the task table **502** is a comprehensive listing of all tasks in a computer cluster such as, for example, the computer cluster **100** of FIG. The task table **502** is typically maintained separately from any given flow-control instance. The task table **502** is regularly updated to add new tasks that are created and to remove tasks that have been completed.

The one or more flow-control instances **518** manage tasks of a managed task type as described with respect to FIG. **2**. Tasks having a same task key, as determined by information from the task table, are grouped into a common reservation. As described above, the task key usually defines a set of characteristics that justify grouping tasks into a same reservation. The set of characteristics can include, for example, a service ID for a requested service with which the task is associated, a customer ID corresponding to a requesting customer, a priority level assigned to the task, and the like.

In a typical embodiment, flow-control performance can be enhanced when, as illustrated, the one or more flow-control instances **518** include more than one flow-control instance to manage tasks of the managed task type. In these embodiments, flow control is distributed across the one or more flow-control instances **518**, with each flow-control instance operating to derive reservations and perform flow control in a similar fashion. It should be appreciated that the one or more flow-control instances **518** need not perform steps of a process such as, for example, the process **400** of FIG. **4**, at identical intervals. Rather, each flow-control instance can execute such a process independently and thus derive reservations at different intervals. The task table **502** operates as a common source of tasks for the one or more flow-control instances **518**. Each of the flow-control instances **518** can be viewed as a cooperative force that results in resources being assigned to reservations derived from the task table **502**. Therefore, when flow control for the managed task type is distributed as illustrated in FIG. **5**, this cooperative effect causes the one or more flow-control instances **518** to exhibit the emergent property.

FIGS. **6-12** illustrate empirical analysis of reservations in greater detail. In a typical embodiment, functionality described with respect to FIGS. **6-12** can be performed as part of step **404** of the process **400** of FIG. **4**. Empirical analysis typically includes generation of a balanced-utilization partitioning scheme such as, for example, the balanced-utilization partitioning scheme **314** of FIG. **3**. In a typical embodiment, the balanced-utilization partitioning scheme utilizes principles of stationary distributions. As used herein, a stationary distribution refers to a set of values that sum to one. As described below, the values can be, for example, percentages of a set of consumable resources. For purposes of this description, the stationary distribution of one represents one-hundred percent of a given set or subset of consumable resources. In that way, as described below, a balance of resource consumption can be precisely maintained.

FIG. **6** illustrates an exemplary balanced-utilization partitioning scheme **600** that may be utilized by a flow-control instance. For example, the balanced-utilization partitioning scheme **600** may be implemented as the balanced-utilization partitioning scheme **314** of FIG. **3**. The balanced-utilization partitioning scheme **600** includes a partition **602**, a partition **604**, a partition **606**, and a partition **608** that each correspond to an active customer (i.e., four active customers). According to the balanced-utilization scheme **600**, all active customers at a given a point in time share a class of consumable resources equally.

FIG. 7 illustrates an implementation 700 of the balanced-utilization partitioning scheme 600 of FIG. 6. The implementation 700 includes a current set of consumable resources 702 that is apportioned among customer-specific reservation sets 706(1), 706(2), 706(3), 706(4), and 706(5) (collectively, customer-specific reservation sets 706). As shown, the current set of consumable resources 702 is apportioned to the customer-specific reservation sets 706 according to a stationary distribution 704. The current set of consumable resources 702 includes N resources.

Each of the customer-specific reservation sets 706 is an aggregation of active reservations for a particular customer (for each of five customers as shown). In particular, the customer-specific reservation set 706(1) includes reservations 711a, 711b, and 711c. The customer-specific reservation set 706(2) includes reservations 711d and 711e. The customer-specific reservation set 706(3) includes reservations 711f, 711g, 711h, and 711i. The customer-specific reservation set 706(4) includes reservation 711j. Finally, the customer-specific reservation set 706(5) includes reservation 711k. For convenient reference, the reservations 711a-k may be referenced collectively as reservations 711.

The stationary distribution 704 functions to distribute customer-specific resource percentages 708(1), 708(2), 708(3), 708(4), and 708(5) to the customer-specific reservation sets 706(1), 706(2), 706(3), 706(4), and 706(5), respectively. For convenient reference, the customer-specific resource percentages 708(1)-708(5) may be referenced collectively as customer-specific resource percentages 708. In compliance with the stationary distribution 704, the customer-specific resource percentages 708 are values that, when summed, equal one. For example, when applying the balanced-utilization scheme 600 of FIG. 6, the current set of consumable resources 702 is apportioned equally. According to this example, each of the customer-specific reservation sets 708 is apportioned one-fifth of the current set of consumable resources 702. As such, each of the customer-specific resource percentages 708 would equal 0.2 in decimal form.

After the stationary distribution 704 is applied, stationary distributions 710(1), 710(2), 710(3), 710(4), and 710(5) are applied (collectively, stationary distributions 710). The stationary distributions 710 apportion the customer-specific resource percentages 708 to individual reservations of the customer-specific reservation sets 706. More particularly, the stationary distributions 710(1), 710(2), 710(3), 710(4), and 710(5) serve to distribute reservation-specific resource percentages 712a-712c, 712d-712e, 712f-712i, 712j, and 712k, respectively. In a typical embodiment, the stationary distributions 710 effect an equal apportionment of the customer-specific resource percentages 708 across each reservation set of the customer-specific reservation sets 706.

More particularly, the stationary distribution 710(1) apportions the customer-specific resource percentage 708(1) to the reservations 711a-711c (i.e., the customer-specific reservation set 706(1)). In this fashion, the reservation-specific resource percentages 712a, 712b, and 712c are distributed to the reservations 711a, 711b, and 711c, respectively. The reservation-specific resource percentages 712a-712c thereby represent proportions of the customer-specific resource percentage 708(1) that, according to the stationary distribution 710(1), collectively sum to one.

The stationary distribution 710(2) apportions the customer-specific resource percentage 708(2) to the reservations 711d-711e (i.e., the customer-specific reservation set 706(2)). In this fashion, the reservation-specific resource percentages 712d and 712e are distributed to the reservations 711d and 711e, respectively. The reservation-specific resource per-

centages 712d and 712e thereby represent proportions of the customer-specific resource percentage 708(2) that, according to the stationary distribution 710(2), collectively sum to one.

The stationary distribution 710(3) apportions the customer-specific resource percentage 708(3) to the reservations 711f-711i (i.e., the customer-specific reservation set 706(3)). In this fashion, the reservation-specific resource percentages 712f, 712g, 712h, and 712i are distributed to the reservations 711f, 711g, 711h, and 711i, respectively. The reservation-specific resource percentages 712f-712i thereby represent proportions of the customer-specific resource percentage 708(3) that, according to the stationary distribution 710(3), collectively sum to one.

The stationary distribution 710(4) apportions the customer-specific resource percentage 708(4) to the reservation 711j (i.e., the customer-specific reservation set 706(4)). In this fashion, the reservation-specific resource percentage 712j is distributed to the reservation 711j. The reservation-specific resource percentage 712j thereby represents a proportion of the customer-specific resource percentage 708(4). Since the customer-specific reservation set 706(4) includes only the reservation 711j, according to the stationary distribution 710(4), the reservation-specific resource percentage 712j will generally equal one.

The stationary distribution 710(5) apportions the customer-specific resource percentage 708(5) to the reservation 711k (i.e., the customer-specific reservation set 706(5)). In this fashion, the reservation-specific resource percentage 712k is distributed to the reservation 711k. The reservation-specific resource percentage 712k thereby represents a proportion of the customer-specific resource percentage 708(5). Since the customer-specific reservation set 706(5) includes only the reservation 711k, according to the stationary distribution 710(5), the reservation-specific resource percentage 712k will generally equal one.

In a typical embodiment, the stationary distributions 710 effect an equal distribution of the customer-specific resource percentages 708 across each reservation set of the customer-specific reservation sets 706. For example, the customer-specific reservation set 706(1) includes three reservations, i.e., the reservations 711a, 711b, and 711c. The reservation-specific resource percentages 712a, 712b, and 712c should thus each equal one-third. The customer-specific reservation set 706(2) includes two reservations, i.e., the reservations 711d and 711e. The reservation-specific resource percentages 712d and 712e should thus each equal one-half. The customer-specific reservation set 706(3) includes four reservations, i.e., the reservations 711f, 711g, 711h, and 711i. The reservation-specific resource percentages 712f, 712g, 712h, and 712i should thus each equal one-fourth. The customer-specific reservation sets 706(4) and 706(5) each include a single reservation, i.e., the reservations 711j and 711k, respectively. Therefore, as described above, the reservation-specific resource percentages 712j and 712k should each equal one.

After the stationary distributions 710 are applied, effective distributions 714a-714k are computed (collectively, effective distributions 714). As explained above, the reservation-specific resource percentages 712 are percentages of the customer-specific resource percentages 708 that should be allocated to the reservations 711. The effective distributions 714 are, in effect, a translation of the reservation-specific resource percentages 712 into percentages of the current set of consumable resources 702. The effective distributions 714a-714k are computed relative to the reservations 711a-711k, respectively.

11

Specifically, each of the effective distributions **714** can be computed as a product of a corresponding reservation-specific resource percentage (from the reservation-specific resource percentages **712**) and a corresponding customer-specific resource percentage (from the customer-specific resource percentages **708**). For example, the effective distribution **714a** can be computed as a product of the reservation-specific resource percentage **712a** and the customer-specific resource percentage **708(1)**. Table 1 below lists exemplary values relative to the example of FIG. 7. It should be appreciated that the effective distributions **714** should sum to one.

TABLE 1

EFFECTIVE DISTRIBUTION	VALUE
Effective distribution 714a	0.0666666666666667
Effective distribution 714b	0.0666666666666667
Effective distribution 714c	0.0666666666666667
Effective distribution 714d	0.1
Effective distribution 714e	0.1
Effective distribution 714f	0.05
Effective distribution 714g	0.05
Effective distribution 714h	0.05
Effective distribution 714i	0.05
Effective distribution 714j	0.2
Effective distribution 714k	0.2

Once the effective distributions **714** have been calculated, in a typical embodiment, flow-control clocking weights **716a-716k** are calculated (collectively, flow-control clocking weights **716**). The flow-control clocking weights **716** are, in effect, a translation of the effective distributions **714** into defined quantities of resources that should be allocated to each of the reservations **711**. As explained in detail below, the flow-control clocking weights **716** can be calculated as products of the effective distributions **714** and a total number of resources in the current set of consumable resources **702** (i.e., **N**).

In particular, the flow-control clocking weight **716a** equals the effective distribution **714a** multiplied by **N**. The flow-control clocking weight **716b** equals the effective distribution **714b** multiplied by **N**. The flow-control clocking weight **716c** equals the effective distribution **714c** multiplied by **N**. The flow-control clocking weight **716d** equals the effective distribution **714d** multiplied by **N**. The flow-control clocking weight **716e** equals the effective distribution **714e** multiplied by **N**. The flow-control clocking weight **716f** equals the effective distribution **714f** multiplied by **N**. The flow-control clocking weight **716g** equals the effective distribution **714g** multiplied by **N**. The flow-control clocking weight **716h** equals the effective distribution **714h** multiplied by **N**. The flow-control clocking weight **716i** equals the effective distribution **714i** multiplied by **N**. The flow-control clocking weight **716j** equals the effective distribution **714j** multiplied by **N**. The flow-control clocking weight **716k** equals the effective distribution **714k** multiplied by **N**.

Each of the flow-control clocking weights **716** is representative of a defined number of resources from the current set of resources **702**. As illustrated, a sum **718** of the flow-control clocking weights **716** equals the total number of consumable resources (i.e., **N**). Therefore, each of the flow-control clocking weights may be expressed in fractional units of resources. As described in greater detail with respect to FIGS. 13-14, the flow-clocking weights **716** enable the balanced-utilization partitioning scheme to be precisely executed on a per reservation basis.

FIG. 8 illustrates an exemplary balanced-utilization partitioning scheme **800** that may be utilized by a flow-control

12

instance. For example, the balanced-utilization partitioning scheme **800** may be implemented as the balanced-utilization partitioning scheme **314** of FIG. 3. The balanced-utilization partitioning scheme **800** includes two super partitions: a high-priority super partition **802** and a regular-priority super partition **804**.

In a typical embodiment, the high-priority super partition **802** includes a subset of a current set of consumable resources that is allocated to eligible active customers. The eligible active customers may include, for example, customers who have elected to pay more for a higher quality of service. In various embodiments, the subset of the current set of consumable resources may be expressed as a percentage, an integer, or in other ways that will be apparent to one of ordinary skill in the art after reviewing the inventive principles described herein. The high-priority super partition **802** is then shared among the active eligible customers as described with respect to the balanced-utilization partitioning scheme **600** of FIG. 6.

In various embodiments, the regular-priority super partition **804** includes all consumable resources of the current set of consumable resources except those consumable resources allocated to the high-priority super partition **802**. In various embodiments, the regular-priority super partition **804** is shared among active customers not eligible for the high-priority super partition **802** in the manner described with respect to FIG. 4. In various other embodiments, the regular-priority super partition **804** is shared among all active customers, including those eligible for the high-priority super partition **802**, in the manner described with respect to FIG. 4.

As described with respect to the balanced-utilization partitioning scheme **214** of FIG. 2A, the balanced-utilization partitioning scheme **800** dynamically changes during execution of the flow-control instance. Correspondingly, in a typical embodiment, a size of the high-priority super partition **802** and a size of the regular-priority super partition **804** can also dynamically change. For example, if fewer than all consumable resources allocated to the high-priority super partition **802** are in use, those unused consumable resources can be allocated to the regular-priority super partition **804** until there are active eligible customers for the high-priority super partition **802** requiring their use.

FIG. 9 illustrates an implementation **900** of the balanced-utilization partitioning scheme **800** of FIG. 8. The implementation **900** includes a regular-priority super partition **902a** and a high-priority super partition **902b**. The regular-priority super partition **902** and the high-priority super partition **902b** generally correspond to the regular-priority super partition **804** and the high-priority super partition **802**, respectively, of FIG. 8. The regular-priority super partition **902a** and the high-priority super partition **902b** each comprise a defined number of consumable resources (at a given time).

In a typical embodiment, the implementation **900** performs an equal distribution **904a** of the regular-priority super partition **902a** across customers **901(1)-901(5)**. In a typical embodiment, the customers **901(1)-901(5)** represent those customers having active reservations (collectively, customers **901**). Therefore, the customers **901(1)-901(5)** are apportioned defined quantities **903(1)-903(5)**, respectively, of resources from the regular-priority super partition **902a**. In a typical embodiment, the defined quantities **903(1)-903(5)**, when summed, should equal a total number of resources contained within the regular-priority super partition **902a**. According to the equal apportionment **904a**, the defined quantities **903(1)-903(5)** represent equal shares of the regular-priority super partition **902a**.

In a typical embodiment, the implementation **900** performs an unequal distribution **904b** of the high-priority super parti-

tion **904b** across active customers who are eligible for a higher quality of service, i.e., the customers **901(2)** and **901(4)**. In a typical embodiment, the customers **901(2)** and **902(4)** are each eligible for predetermined resource quantities **905(2)** and **905(4)**, respectively, of the high-priority super partition **904b**. According to the unequal apportionment **904b**, the predetermined resource quantities **905(2)** and **905(4)** are not necessarily equal, although equality is possible.

As a result of the equal apportionment **904a** and the unequal apportionment **904b**, the customers **901(1)**-**901(5)** have resource apportionments **906(1)**-**906(5)**, respectively. The resource apportionments **906(1)**-**906(5)** are typically sums of quantities apportioned via the equal apportionment **904a** and quantities apportioned via the unequal apportionment **904b**. In particular, as shown, the customers **901(1)**, **901(3)**, and **901(5)** are not eligible for the high-priority super partition **902b**. Therefore, the resource apportionments **906(1)**, **906(3)**, and **906(5)** equal the defined quantities **903(1)**, **903(3)**, and **903(5)**, respectively. Since the customer **901(2)** is eligible for the high-priority super partition **902b**, the resource apportionment **906(2)** equals a sum of the defined quantity **903(2)** and the predetermined resource quantity **905(2)**. In like manner, since the customer **901(4)** is also eligible for the high-priority super partition **902b**, the resource apportionment **906(4)** equals a sum of the defined quantity **903(4)** and the predetermined resource quantity **905(4)**.

Total apportioned resources **906** are an aggregation of the resource apportionments **906(1)**-**906(5)**. As described with respect to FIG. 8, as an optimization, resources of the high-priority super partition **902b** that are not being utilized are allocated to the regular-priority super partition **902a**. Therefore, in some embodiments, the total apportioned resources **906** will include a quantity of resources equal in number to a combined total number of resources in the regular-priority super partition **902a** and the high-priority super partition **902b**. In various other embodiments, it should be appreciated that, by design, this may not be the case. For purposes of this example, the quantity of resources in the total apportioned resources **906** may be considered to be **P**.

Once the resource apportionments **906(1)**-**906(5)** are computed, a number-to-percentage distribution **908** can be applied. As described above, the resource apportionments **906(1)**-**906(5)** represent quantities of resources. The number-to-percentage distribution **908** operates to convert the resource apportionments **906(1)**-**906(5)** to percentages that are distributed to the customer-specific reservation sets **910(1)**-**910(5)** (collectively, customer-specific reservation sets **910**), with each reservation set constituting an aggregation of customer reservations.

In particular, the customer-specific reservation sets **910(1)**-**910(5)** are aggregations of active reservations for the customers **901(1)**-**901(5)**, respectively. The customer-specific reservation set **910(1)** includes reservations **915a**, **915b**, and **915c**. The customer-specific reservation set **910(2)** includes reservations **915d** and **915e**. The customer-specific reservation set **910(3)** includes reservations **915f**, **915g**, **915h**, and **915i**. The customer-specific reservation set **910(4)** includes reservation **915j**. Finally, the customer-specific reservation set **910(5)** includes reservation **915k**. For convenient reference, the reservations **915a-k** may be referenced collectively as reservations **915**.

The number-to-percentage distribution **908** functions to distribute customer-specific resource percentages **912(1)**, **912(2)**, **912(3)**, **912(4)**, and **912(5)** to the customer-specific reservation sets **910(1)**, **910(2)**, **910(3)**, **910(4)**, and **910(5)**, respectively. For convenient reference, the customer-specific resource percentages **912(1)**-**912(5)** may be referenced col-

lectively as customer-specific resource percentages **912**. For example, the customer-specific resource percentages **912(1)**-**912(5)** can equal the resource apportionments **906(1)**-**906(5)**, respectively, divided by a total number of resources in the total apportioned resources **906**. For example, if the resource apportionment **906(1)** included five resources and the total apportioned resources **906** included one-hundred resources, the customer-specific resource percentage **912(1)** would equal 0.05 in decimal form.

After the number-to-percentage distribution **908** is applied, stationary distributions **914(1)**, **914(2)**, **914(3)**, **914(4)**, and **914(5)** are applied (collectively, stationary distributions **914**). The stationary distributions **914** apportion the customer-specific resource percentages **912** to individual reservations of the customer-specific reservation sets **910**. Specifically, the stationary distributions **914(1)**, **914(2)**, **914(3)**, **914(4)**, and **914(5)** serve to distribute reservation-specific resource percentages **916a**-**916c**, **916d**-**916e**, **916f**-**916i**, **916j**, and **916k**, respectively. In a typical embodiment, the stationary distributions **914** effect an equal apportionment of the customer-specific resource percentages **912** across each reservation set of the customer-specific reservation sets **910**.

More particularly, the stationary distribution **914(1)** apportions the customer-specific resource percentage **912(1)** to the reservations **915a**-**915c** (i.e., the customer-specific reservation set **910(1)**). In this fashion, the reservation-specific resource percentages **916a**, **916b**, and **916c** are distributed to the reservations **915a**, **915b**, and **915c**, respectively. The reservation-specific resource percentages **916a**-**916c** thereby represent proportions of the customer-specific resource percentage **912(1)** that, according to the stationary distribution **914(1)**, collectively sum to one.

The stationary distribution **914(2)** apportions the customer-specific resource percentage **912(2)** to the reservations **915d**-**915e** (i.e., the customer-specific reservation set **910(2)**). In this fashion, the reservation-specific resource percentages **916d** and **916e** are distributed to the reservations **915d** and **915e**, respectively. The reservation-specific resource percentages **916d** and **916e** thereby represent proportions of the customer-specific resource percentage **912(2)** that, according to the stationary distribution **914(2)**, collectively sum to one.

The stationary distribution **914(3)** apportions the customer-specific resource percentage **912(3)** to the reservations **915f**-**915i** (i.e., the customer-specific reservation set **910(3)**). In this fashion, the reservation-specific resource percentages **916f**, **916g**, **916h**, and **916i** are distributed to the reservations **915f**, **915g**, **915h**, and **915i**, respectively. The reservation-specific resource percentages **916f**-**916i** thereby represent proportions of the customer-specific resource percentage **912(3)** that, according to the stationary distribution **914(3)**, collectively sum to one.

The stationary distribution **914(4)** apportions the customer-specific resource percentage **912(4)** to the reservation **915j** (i.e., the customer-specific reservation set **910(4)**). In this fashion, the reservation-specific resource percentage **916j** is distributed to the reservation **915j**. The reservation-specific resource percentage **916j** thereby represents a proportion of the customer-specific resource percentage **912(4)**. Since the customer-specific reservation set **910(4)** includes only the reservation **915j**, according to the stationary distribution **914(4)**, the reservation-specific resource percentage **916j** will generally equal one.

The stationary distribution **914(5)** apportions the customer-specific resource percentage **912(5)** to the reservation **915k** (i.e., the customer-specific reservation set **910(5)**). In this fashion, the reservation-specific resource percentage **916k** is distributed to the reservation **915k**. The reservation-

15

specific resource percentage **916k** thereby represents a proportion of the customer-specific resource percentage **912(5)**. Since the customer-specific reservation set **910(5)** includes only the reservation **915k**, according to the stationary distribution **914(5)**, the reservation-specific resource percentage **916k** will generally equal one.

In a typical embodiment, the stationary distributions **914** effect an equal distribution of the customer-specific resource percentages **912** across each reservation set of the customer-specific reservation sets **910**. For example, the customer-specific reservation set **910(1)** includes three reservations, i.e., the reservations **915a**, **915b**, and **915c**. The reservation-specific resource percentages **916a**, **916b**, and **916c** should thus each equal one-third. The customer-specific reservation set **910(2)** includes two reservations, i.e., the reservations **915d** and **915e**. The reservation-specific resource percentages **916d** and **916e** should thus each equal one-half. The customer-specific reservation set **910(3)** includes four reservations, i.e., the reservations **915f**, **915g**, **915h**, and **915i**. The reservation-specific resource percentages **916f**, **916g**, **916h**, and **916i** should thus each equal one-fourth. The customer-specific reservation sets **910(4)** and **910(5)** each include a single reservation, i.e., the reservations **915j** and **915k**, respectively. Therefore, as described above, the reservation-specific resource percentages **916j** and **916k** should each equal one.

After the stationary distributions **914** are applied, effective distributions **918a-918k** are computed (collectively, effective distributions **918**). As explained above, the reservation-specific resource percentages **916** are percentages of the customer-specific resource percentages **912** that should be allocated to the reservations **915**. The effective distributions **918** are, in effect, a translation of the reservation-specific resource percentages **916** into percentages of the total apportioned resources **906**. The effective distributions **918a-918k** are computed relative to the reservations **915a-915k**, respectively.

Specifically, each of the effective distributions **918** can be computed as a product of a corresponding reservation-specific resource percentage (from the reservation-specific resource percentages **916**) and a corresponding customer-specific resource percentage (from the customer-specific resource percentages **912**). For example, the effective distribution **918a** can be computed as a product of the reservation-specific resource percentage **916a** and the customer-specific resource percentage **912(1)**.

Once the effective distributions **918** have been calculated, in a typical embodiment, flow-control clocking weights **920a-920k** are calculated (collectively, flow-control clocking weights **920**). The flow-control clocking weights **920** are, in effect, a translation of the effective distributions **918** into defined quantities of resources that should be allocated to each of the reservations **915**. As explained in detail below, the flow-control clocking weights **920** can be calculated as products of the effective distributions **918** and a total number of apportioned resources in the total apportioned resources **906** (i.e., **P**).

In particular, the flow-control clocking weight **920a** equals the effective distribution **918a** multiplied by **P**. The flow-control clocking weight **920b** equals the effective distribution **918b** multiplied by **P**. The flow-control clocking weight **920c** equals the effective distribution **918c** multiplied by **P**. The flow-control clocking weight **920d** equals the effective distribution **918d** multiplied by **P**. The flow-control clocking weight **920e** equals the effective distribution **918e** multiplied by **P**. The flow-control clocking weight **920f** equals the effective distribution **918f** multiplied by **P**. The flow-control clock-

16

ing weight **920g** equals the effective distribution **918g** multiplied by **P**. The flow-control clocking weight **920h** equals the effective distribution **918h** multiplied by **P**. The flow-control clocking weight **920i** equals the effective distribution **918i** multiplied by **P**. The flow-control clocking weight **920j** equals the effective distribution **918j** multiplied by **P**. The flow-control clocking weight **920k** equals the effective distribution **918k** multiplied by **P**.

Each of the flow-control clocking weights **920** is representative of a defined number of resources from the total apportioned resources **906**. As illustrated, a sum **922** of the flow-control clocking weights **920** equals the total number of apportioned resources (i.e., **P**). Therefore, each of the flow-control clocking weights may be expressed in fractional units of resources. As described in greater detail with respect to FIGS. 13-14, the flow-clocking weights **920** enable the balanced-utilization partitioning scheme to be precisely executed on a per reservation basis.

FIG. 10 illustrates a balanced-utilization partitioning scheme **1000** that may be utilized by a flow-control instance such as, for example, the flow-control instance **200** of FIG. 2A. For example, the balanced-utilization partitioning scheme **1000** may be implemented as the balanced-utilization partitioning scheme **214** of FIG. 2A. The balanced-utilization partitioning scheme **1000** includes a high-priority super partition **1002** and a regular-priority super partition **1004**. The high-priority super partition **1002** operates as described with respect to the high-priority super partition **802** of FIG. 8 and the high-priority super partition **902b** of FIG. 9. The regular-priority super partition **1004** operates as described with respect to the regular-priority super partition **804** of FIG. 8 and the regular-priority super partition **902a** of FIG. 9. However, differently from FIGS. 8-9, the balanced-utilization scheme **1000** additionally includes a low-priority super partition **1004a** that is a subset of the regular-priority super partition **1004**.

The low-priority super partition **1004a** is a partition of dynamically variable size and includes consumable resources of the regular-priority super partition **1004** that are not in use at a given time. In a typical embodiment, the low-priority super partition **1004a** allows a designated class of customers to pay a much lower price for service in exchange for a much lower quality of service. In addition, the low-priority super partition **1004a** helps to maximize overall utilization of the current set of consumable resources. Typically, the low-priority super partition **1004a** is active only when the regular-priority super partition **1004** is not being fully utilized by eligible active customers. Whenever there are sufficient eligible active customers to fully utilize the regular-priority super partition **1004**, the low-priority super partition **1004a** is typically inactive.

In various embodiments, reservations can be associated with priority levels selected from a plurality of priority levels (e.g., priority levels one to nine). In a typical embodiment, the plurality of priority levels vary from low priority (e.g., level one) to high priority (e.g., level nine). As described in more detail below, a priority level associated with a given reservation can be used to configurably vary a computation of effective distributions.

FIG. 11 illustrates an implementation **1100** of a balanced-utilization partitioning scheme that utilizes priorities. The implementation **1100** begins by apportioning a set of consumable resources among customer-specific reservation sets **1104(1)**, **1104(2)**, **1104(3)**, **1104(4)**, and **1104(5)** (collectively, customer-specific reservation sets **1104**). As shown,

the set of consumable resources is apportioned to the customer-specific reservation sets **1104** according to a distribution **1102**.

Each of the customer-specific reservation sets **1104** is an aggregation of active reservations for a particular customer (for each of five customers as shown). In particular, the customer-specific reservation set **1104(1)** includes reservations **1114a**, **1114b**, and **1114c**. The customer-specific reservation set **1104(2)** includes reservations **1114d** and **1114e**. The customer-specific reservation set **1104(3)** includes reservations **1114f**, **1114g**, **1114h**, and **1114i**. The customer-specific reservation set **1104(4)** includes reservation **1114j**. Finally, the customer-specific reservation set **1104(5)** includes reservation **1114k**. For convenient reference, the reservations **1114a-1114k** may be referenced collectively as reservations **1114**. The reservations **114** each have a priority level associated therewith. In a typical embodiment, the priority level may be selected from a plurality of priority levels of increasing priority (e.g., one to nine).

In various embodiments, the distribution **1102** may be applied as described, for example, with respect to the number-to-percentage distribution **908** of FIG. **9** or the stationary distribution **704** of FIG. **7**. In that way, the distribution **1102** may apply a balanced partition scheme as described with respect to FIG. **7** or apply a high-priority super partition and a regular-priority super partition as described with respect to FIG. **9**. Thus, the distribution **1102** distributes customer-specific resource percentages **1103(1)-1103(5)** to the customer-specific reservation sets **1104(1)-1104(5)**, respectively. For convenient reference, the customer-specific resource percentages **1103(1)-1103(5)** may be referenced collectively as resource percentages **1103**.

In a typical embodiment, after the distribution **1102** has been applied, stationary distributions **1106(1)-1106(5)** are applied (collectively, stationary distributions **1106**). The stationary distributions **1106** apportion the customer-specific resource percentages **1103** to priority groupings **1108a-1108g** of the customer-specific reservation sets **1104**. In a typical embodiment, a priority grouping exists for each priority level at which there is an active reservation for a given customer. In particular, the stationary distributions **1106(1)**, **1106(2)**, **1106(3)**, **1106(4)**, and **1106(5)** serve to distribute priority-specific resource percentages **1110a-1110b**, **1110c**, **1110d-1110e**, **1110f**, and **1110g** respectively. In a typical embodiment, priority groupings of the priority groupings **1108** that represent higher priority levels are awarded greater proportions of the customer-specific resource percentages **1103** than those priority groupings representing lower priority levels. An example of how the priority-specific resource percentages can be computed will be described with respect to FIG. **12**.

The stationary distribution **1106(1)** apportions the customer-specific resource percentage **1103(1)** to the priority groupings **1108a** and **1108b**. The priority groupings **1108a** and **1108b** indicate that the customer-specific reservation set **1104(1)** includes one or more reservations at priority-level one and priority-level two, respectively (i.e. two priority groupings). The priority-specific resource percentages **1110a** and **1110b** thereby represent proportions of the customer-specific resource percentage **1103(1)** that are apportioned to the priority groupings **1108a** and **1108b**, respectively. According to the stationary distribution **1106(1)**, the priority-specific resource percentages **1110a** and **1110b** collectively sum to one.

The stationary distribution **1106(2)** apportions the customer-specific resource percentage **1103(2)** to the priority grouping **1108c**. The priority grouping **1108c** indicates that

the customer-specific reservation set **1104(2)** includes one or more reservations at priority-level three (i.e., one priority grouping). The priority-specific resource percentage **1110c** thereby represents a proportion of the customer-specific resource percentage **1103(2)** that is apportioned to the priority grouping **1108c**. Since the customer-specific reservation set **1104(2)** only includes one priority grouping, i.e., the priority grouping **1108c**, according to the stationary distribution **1106(2)**, the priority-specific resource percentage **1110c** should equal one.

The stationary distribution **1106(3)** apportions the customer-specific resource percentage **1103(3)** to the priority groupings **1108d** and **1108e**. The priority groupings **1108d** and **1108e** indicate that the customer-specific reservation set **1104(3)** includes one or more reservations at priority-level one and priority-level two, respectively (i.e. two priority groupings). The priority-specific resource percentages **1110d** and **1110e** thereby represent proportions of the customer-specific resource percentage **1103(3)** that are apportioned to the priority groupings **1108d** and **1108e**, respectively. According to the stationary distribution **1106(3)**, the priority-specific resource percentages **1110d** and **1110e** collectively sum to one.

The stationary distribution **1106(4)** apportions the customer-specific resource percentage **1103(4)** to the priority grouping **1108f**. The priority grouping **1108f** indicates that the customer-specific reservation set **1104(4)** includes one or more reservations at priority-level one (i.e., one priority grouping). The priority-specific resource percentage **1110f** thereby represents a proportion of the customer-specific resource percentage **1103(4)** that is apportioned to the priority grouping **1108f**. Since the customer-specific reservation set **1104(4)** only includes one priority grouping, i.e., the priority grouping **1108f**, according to the stationary distribution **1106(4)**, the priority-specific resource percentage **1110f** should equal one.

The stationary distribution **1106(5)** apportions the customer-specific resource percentage **1103(5)** to the priority grouping **1108g**. The priority grouping **1108g** indicates that the customer-specific reservation set **1104(5)** includes one or more reservations at priority-level one (i.e., one priority grouping). The priority-specific resource percentage **1110g** thereby represents a proportion of the customer-specific resource percentage **1103(5)** that is apportioned to the priority grouping **1108g**. Since the customer-specific reservation set **1104(5)** only includes one priority grouping, i.e., the priority grouping **1108g**, according to the stationary distribution **1106(5)**, the priority-specific resource percentage **1110g** should equal one.

After the stationary distributions **1106** are applied, stationary distributions **1112(1)**, **1112(2)**, **1112(3)**, **1112(4)**, **1112(5)**, **1112(6)**, and **1112(7)** are applied (collectively, stationary distributions **1112**). The stationary distributions **1112** apportion the priority-specific resource percentages **1110** to individual reservations of the priority groupings **1108**. Specifically, the stationary distributions **1112(1)**, **1112(2)**, **1112(3)**, **1112(4)**, **1112(5)**, **1112(6)**, and **1112(7)** serve to distribute reservation-specific resource percentages **1116a**, **1116b-1116c**, **1116d-1116e**, **1116f-1116g**, **1116h-1116i**, **1116j**, and **1116k** respectively. In a typical embodiment, the stationary distributions **1112** effect an equal distribution of the priority-specific resource percentages **1110** across each priority grouping of the priority groupings **1108**.

More particularly, the stationary distribution **1112(1)** apportions the priority-specific resource percentage **1110a** to the reservation **1114a** (i.e., the priority grouping **1108a**). In this fashion, the reservation-specific resource percentage

1116a is distributed to the reservation 1114a. The reservation-specific resource percentage 1116a thereby represents a proportion of the priority-specific resource percentage 1110a. Since the priority grouping 1108a only includes one reservation (i.e., the reservation 1114a), according to the stationary distribution 1112(1), the reservation-specific resource percentage 1116a should equal one.

The stationary distribution 1112(2) apportions the priority-specific resource percentage 1110b to the reservations 1114b-1114c (i.e., the priority grouping 1108b). In this fashion, the reservation-specific resource percentages 1116b and 1116c are distributed to the reservations 1114b and 1114c, respectively. The reservation-specific resource percentages 1116b and 1116c thereby represent proportions of the priority-specific resource percentage 1110b that, according to the stationary distribution 1112(2), collectively sum to one.

The stationary distribution 1112(3) apportions the priority-specific resource percentage 1110c to the reservations 1114d-1114e (i.e., the priority grouping 1108c). In this fashion, the reservation-specific resource percentages 1116d and 1116e are distributed to the reservations 1114d and 1114e, respectively. The reservation-specific resource percentages 1116d and 1116e thereby represent proportions of the priority-specific resource percentage 1110c that, according to the stationary distribution 1112(3), collectively sum to one.

The stationary distribution 1112(4) apportions the priority-specific resource percentage 1110d to the reservations 1114f-1114g (i.e., the priority grouping 1108d). In this fashion, the reservation-specific resource percentages 1116f and 1116g are distributed to the reservations 1114f and 1114g, respectively. The reservation-specific resource percentages 1116f and 1116g thereby represent proportions of the priority-specific resource percentage 1110d that, according to the stationary distribution 1112(4), collectively sum to one.

The stationary distribution 1112(5) apportions the priority-specific resource percentage 1110e to the reservations 1114h-1114i (i.e., the priority grouping 1108e). In this fashion, the reservation-specific resource percentages 1116h and 1116i are distributed to the reservations 1114h and 1114i, respectively. The reservation-specific resource percentages 1116h and 1116i thereby represent proportions of the priority-specific resource percentage 1110e that, according to the stationary distribution 1112(5), collectively sum to one.

The stationary distribution 1112(6) apportions the priority-specific resource percentage 1110f to the reservation 1114j (i.e., the priority grouping 1108f). In this fashion, the reservation-specific resource percentage 1116j is distributed to the reservation 1114j. The reservation-specific resource percentage 1116j thereby represents a proportion of the priority-specific resource percentage 1110f. Since the priority grouping 1108f only includes one reservation (i.e., the reservation 1114j), according to the stationary distribution 1112(6), the reservation-specific resource percentage 1116j should equal one.

The stationary distribution 1112(7) apportions the priority-specific resource percentage 1110g to the reservation 1114k (i.e., the priority grouping 1108g). In this fashion, the reservation-specific resource percentage 1116k is distributed to the reservation 1114k. The reservation-specific resource percentage 1116k thereby represents a proportion of the priority-specific resource percentage 1110g. Since the priority grouping 1108g only includes one reservation (i.e., the reservation 1114k), according to the stationary distribution 1112(7), the reservation-specific resource percentage 1116k should equal one.

In a typical embodiment, the stationary distributions 1112 effect an equal distribution of the priority-specific resource

percentages 1110 across each priority grouping of the priority groupings 1108. For example, the priority groupings 1108a, 1110f, and 1110g each include one reservation. Therefore, the reservation-specific resource percentages 1116a, 1116j, and 1116j should each equal one. By way of further example, the priority groupings 1108b, 1108c, 1108d, and 1108e each include two reservations. Therefore, the reservation-specific resource percentages 1116b-1116i should each equal one-half.

After the stationary distributions 1112 are applied, effective distributions 1118a-1118k are computed (collectively, effective distributions 1118). As explained above, the reservation-specific resource percentages 1116 are percentages of the priority-specific resource percentages 1110 that should be allocated to the reservations 1114. The effective distributions 1118 are, in effect, a translation of the reservation-specific resource percentages 1116 into percentages of the current set of consumable resources. The effective distributions 1118a-1118k are computed relative to the reservations 1114a-1114k, respectively.

The distribution 1102 may be considered an outer stationary distribution since it is across customers and thus “outside” any one customer. The stationary distributions 1106 and the stationary distributions 1108 may be considered inner stationary distributions as since they are computed with respect to particular customers and are thus “inside” particular customers. Each of the effective distributions 1118 represents a proportion of the current set of consumable resources that is being apportioned to a given reservation of the reservations 1114. Each of the effective distributions 1118 can be computed as a product of the outer stationary distribution and each inner stationary distribution.

Stated differently, each of the effective distributions 1118 can be computed as a product of a corresponding reservation-specific resource percentage (from the reservation-specific resource percentages 1116), a corresponding priority-specific resource percentage (from the priority-specific resource percentages 1110), and a corresponding customer-specific resource percentage (from the customer-specific resource percentages 1103). For example, the effective distribution 1118a can be computed as a product of the reservation-specific resource percentage 1116a, the priority-specific resource percentage 1110a, and the customer-specific resource percentage 1103(1). It should be appreciated that the effective distributions 1118 should sum to one.

Once the effective distributions 1118 have been calculated, in a typical embodiment, flow-control clocking weights 1120a-1120k are calculated (collectively, flow-control clocking weights 1120). The flow-control clocking weights 1120 are, in effect, a translation of the effective distributions 1118 into defined quantities of resources that should be allocated to each of the reservations 1114. As explained in detail below, the flow-control clocking weights 1120 can be calculated as products of the effective distributions 1118 and a total number of resources in the current set of consumable resources (i.e., N).

In particular, the flow-control clocking weight 1120a equals the effective distribution 1118a multiplied by N. The flow-control clocking weight 1120b equals the effective distribution 1118b multiplied by N. The flow-control clocking weight 1120c equals the effective distribution 1118c multiplied by N. The flow-control clocking weight 1120d equals the effective distribution 1118d multiplied by N. The flow-control clocking weight 1120e equals the effective distribution 1118e multiplied by N. The flow-control clocking weight 1120f equals the effective distribution 1118f multiplied by N. The flow-control clocking weight 1120g equals the effective

21

distribution **1118g** multiplied by **N**. The flow-control clocking weight **1120h** equals the effective distribution **1118h** multiplied by **N**. The flow-control clocking weight **1120i** equals the effective distribution **1118i** multiplied by **N**. The flow-control clocking weight **1120j** equals the effective distribution **1118j** multiplied by **N**. The flow-control clocking weight **1120k** equals the effective distribution **1118k** multiplied by **N**.

Each of the flow-control clocking weights **1120** is representative of a defined number of resources from the current set of consumable resources. As illustrated, a sum **1122** of the flow-control clocking weights **1120** equals the total number of consumable resources (i.e., **N**). Therefore, each of the flow-control clocking weights may be expressed in fractional units of resources. As described in greater detail with respect to FIGS. **13-14**, the flow-clocking weights **1120** enable the balanced-utilization partitioning scheme to be precisely executed on a per reservation basis.

After reviewing the inventive principles contained herein, one of ordinary skill in the art will appreciate that outer stationary distributions and inner stationary distributions may be established differently than is described with respect to FIG. **11**. For example, any number of inner stationary distributions may be established to further sub-divide a given percentage of consumable resources. In each case, effective distributions may be computed as products of an outer stationary distribution and each inner stationary distribution.

FIG. **12** illustrates a process **1200** for calculating an inner stationary distribution for a customer based on priorities. In a typical embodiment, the process **1200** is performed for each active customer as described, for example, with respect to FIG. **11**. The process **1200** can be performed by a computer cluster such as, for example, the computer cluster **100** of FIG. **1**. The process **1200** begins at step **1202**. At step **1202**, the customer's priority groupings are identified. In a typical embodiment, a priority grouping is identified for each priority level of a plurality of priority levels at which the customer has an active reservation. The plurality of priority levels may be enumerated from 1 to **n**. From step **1202**, the process **1200** proceeds to step **1204**.

At step **1204**, a relative size between priorities is computed for each priority grouping. The relative size represents a proportion of a given set of resources that should be apportioned a given priority level. For example, in one embodiment, the relative size can be calculated using an exponential function such as, for example, the function below, where **p** represents a priority level of the priority grouping and **F(p)** represents the relative size:

$$F(p)=1.25^{p-1}$$

From step **1204**, the process **1200** proceeds to step **1206**.

At step **1206**, a sum of all relative sizes is calculated. From step **1206**, the process **1200** proceeds to step **1208**. At step **1208**, for each priority grouping, a normalized relative size is calculated based on which priority groupings exist. For example, for each priority grouping, the normalized relative size can equal the relative size divided by the sum of all relative sizes. It should be appreciated that a sum of all normalized relative sizes for the customer should equal one (i.e., a stationary distribution). After step **1208**, the process **1200** ends.

FIG. **13** illustrates clocking functionality **1300** of a flow-control instance **1328**. The clocking functionality **1300** may be performed, for example, as part of step **406** of the process **400** of FIG. **4**. In a typical embodiment, the flow-control instance **1328** operates as described with respect to the flow-control instance **300** of FIG. **3**. More particularly, however, the flow-control instance **1328** maintains reservations **1316**

22

(**1**), **1316(2)**, and **1316(3)**. The reservations **1316(1)**, **1316(2)**, and **1316(3)** include tasks **1320(1)**, **1320(2)**, and **1320(3)**, respectively, flow-control clocking weights **1322(1)**, **1322(2)**, and **1322(3)**, respectively, and wait-times **1324(1)**, **1324(2)**, and **1324(3)**, respectively.

For convenient reference, the reservations **1316(1)**, **1316(2)**, and **1316(3)**, the flow-control clocking weights **1322(1)**, **1322(2)**, and **1322(3)**, and the wait-times **1324(1)**, **1324(2)**, and **1324(3)** may be referenced collectively as a current set of active reservations **1316**, clocking weights **1322**, and wait-times **1324**, respectively. For illustrative purposes, the current set of active reservations **1316** is shown to include the reservations **1316(1)**, **1316(2)**, and **1316(3)**, which reservations may be considered to represent an exemplary snapshot-in-time of the current set of active reservations **1316**. As described in more detail below, the flow-control instance **1328** maintains a priority queue **1318** of the current set of active reservations **1316**.

The tasks **1320(1)**, **1320(2)**, and **1320(3)** are each a derived grouping of one or more tasks having, for example, a same task key as described above. In a typical embodiment, the clocking weights **1322** enable precise execution of a balanced-utilization partitioning scheme such as, for example, the balanced-utilization partitioning schemes described with respect to FIGS. **6-12**. The flow-control clocking weights **1322(1)**, **1322(2)**, and **1322(3)** define a proportion of a current set of consumable resources that should be allocated to the reservations **1316(1)**, **1316(2)**, and **1316(3)**, respectively. As described with respect to FIG. **14**, each of the flow-control clocking weights **1316** is computed as a product of an effective distribution and a number of resources in the current set of consumable resources. Therefore, the clocking weights **1322(1)**, **1322(2)**, and **1322(3)** may be expressed in fractional units of resources. It should be appreciated that the clocking weights **1322(1)**, **1322(2)**, and **1322(3)** change dynamically responsive to changes in, for example, the current set of active reservations **1316**.

The wait-times **1324** are an up-to-date measure of how needy the reservations **1316** are for the class of consumable resources and are updated by the flow-control instance **1328**. The wait-times **1324** may be expressed in fractional units of resource-time (e.g., resource-seconds when wall-time is used). The flow-control instance **1328** maintains a virtual clock **1326** that, every clocking cycle, initiates a virtual pendulum for each reservation in the current set of active reservations **1316**. In various embodiments, a clocking cycle can be, for example, a configurable unit of wall time (e.g., seconds or milliseconds) or a configurable number of CPU cycles. Other alternatives that can be considered a clocking cycle will be apparent to one of ordinary skill in the art after reviewing the inventive principles described herein.

As described in greater detail with respect to FIGS. **15-16**, the wait-times **1324(1)**, **1324(2)**, and **1324(3)** are each given initial values upon creation of the reservations **1316(1)**, **1316(2)**, and **1316(3)**, respectively. The initiation of the virtual pendulums results in the reservations **1316** being "clocked" each clocking cycle according to the clocking weights **1322**. It should be noted that each reservation in the current set of active reservations **1316**, by virtue of its inclusion, includes at least one task that has not been processed to completion. Therefore, in various embodiments in which lower wait-times represent greater neediness, the wait-times **1324** are decremented each clocking cycle by a factor of a corresponding clocking weight of the clocking weights **1322** (i.e., clocked to move the virtual pendulum in a negative direction). An example of how the wait-times **1324** can be decremented will be described with respect to FIG. **14**.

When a task of a given reservation in the current set of reservations **1316** is finished accessing a given consumable resource, a corresponding wait-time of the wait-times **1324** is incremented by a factor of a corresponding clocking weight of the clocking weights **1322** (i.e., clocked to move the virtual pendulum in a positive direction). For example, the corresponding wait-time can be incremented by a product of the corresponding clocking weight and a number of clocking units that the task has accessed the given consumable resource. In particular, in embodiments in which wall time is used to define the clocking cycle, the number of clocking units may be expressed in seconds with precision, for example, to a nearest millisecond. Therefore, if the task has accessed the given consumable resource for 12.001 seconds, the corresponding wait-time would be incremented by a product of 12.001 seconds and the corresponding clocking weight (i.e., units of resource-seconds).

The flow-control instance **1328** maintains the priority queue **1318** based on the current set of active reservations **1316**. In a typical embodiment, the priority queue **1318** is sorted by the wait-times **1324**, where a lowest wait-time indicates a most-needy reservation and a highest wait-time indicates a least-needy reservation. Each clocking cycle, as the virtual pendulums are clocked negatively or positively in the manner described above, the flow-control instance **1328** re-sorts the priority queue **1318** based on updated values for the wait-times **1324**. As described in greater detail below with respect to the ensuing Figures, the flow-control instance **1328** assigns consumable resources of the current set of consumable resources to reservations of the current set of active reservations **1316** based on relative neediness. In a typical embodiment, the clocking weights **1322** and the wait-times **1324** may be maintained in a double-precision floating-point format according to the IEEE 754 standard.

FIG. **14** illustrates a process **1400** for decrementing a wait-time for each reservation of a set of active reservations. For example, the process **1400** may be performed as part of performing clocking as described with respect to FIG. **13**. The process **1400** may be performed by a computer cluster such as, for example, the computer cluster **100** of FIG. **1**. The process **1400** is typically performed each clocking cycle. The process **1400** begins at step **1402**.

At step **1402**, a flow-control clocking weight is accessed. In a typical embodiment, the flow-control clocking weight is computed as part of generating a balanced-utilization partitioning scheme as described with respect to FIGS. **7**, **9**, and **11**. As described above, the flow-control clocking weight equals the reservation's effective distribution multiplied by a total number of resources in a current set of consumable resources. In that way, the flow-control clocking weight represents a number of resources and can be expressed in fractional units of resources. From step **1402**, the process **1400** proceeds to step **1404**. At step **1404**, for each resource, resource utilization during the clocking cycle is determined. For example, the computer cluster tracks how many clocking units that each resource has been utilized. The resource utilization can be expressed in fractional units of resource-time. For example, in embodiments in which wall time is utilized, the resource utilization can be expressed in fractional units of resource-seconds to a nearest millisecond. From step **1404**, the process **1400** proceeds to step **1406**.

At step **1406**, resource utilizations are summed across the current set of consumable resources. From step **1406**, the process **1500** proceeds to step **1408**. At step **1408**, a maximum theoretical utilization of the current set of consumable resources is determined. The maximum theoretical utilization corresponds to a number of clocking units that, collectively,

the current set of consumable resources could have been utilized during the clocking cycle. In a typical embodiment, the maximum theoretical utilization equals the total number of resources multiplied by a length of the clocking cycle (e.g., seconds if wall time is utilized). Therefore, the maximum theoretical utilization also has units of resource-time (e.g., resource-seconds if wall time is utilized). From step **1408**, the process **1400** proceeds to step **1410**.

At step **1410**, an actual average capacity of the current set of consumable resources is computed. The actual average capacity can equal, for example, the summed resource utilizations divided by the maximum theoretical utilization. Consequently, the actual average capacity can be represented as a percentage of the maximum theoretical utilization. From step **1410**, the process **1400** proceeds to step **1412**. At step **1412**, the computer cluster computes, for each reservation, a product of the flow-control clocking weight (units of resources), the actual average capacity, and the length of the clocking cycle (time, e.g., in seconds). Therefore, the product can be expressed in resource-seconds. From step **1412**, the process **1400** proceeds to step **1414**. At step **1414**, for each reservation, the wait-time is decremented by the product computed at step **1412**. After step **1414**, the process **1400** ends.

FIG. **15** illustrates a collection **1500** of interactive processes that may be executed by a flow-control instance. The flow-control instance manages a dynamically changing set of consumable resources relative to a managed task type requiring utilization of a class of consumable resources. In a typical embodiment, the collection **1500** is used to control sharing of the set of consumable resources among a dynamically changing set of active customers. The collection **1500** may be executed by a computer cluster such as, for example, the computer cluster **100** of FIG. **1**.

The collection **1500** includes a customer-needs process **1502**, a reservations process **1504**, a reservation-accounting process **1506**, and a resource-allocation process **1508**. In a typical embodiment, each process in the collection **1500** is executed in parallel by one or more computers of the computer cluster. Upon instantiation of the flow-control instance, no process of the collection **1500** is typically in execution. Execution of the collection **1500** typically begins with a first execution of the customer-needs process **1502**. The customer-needs process **1500** begins at step **1502(1)**.

At step **1502(1)**, a customer-specific task group is created. In a typical embodiment, the customer-specific task group is created responsive to a request for service being received from a customer. Creation of the customer-specific task group typically involves grouping tasks, for example, by task key as described above with respect to FIGS. **3-5**. It should be appreciated that the customer-specific task group might include a single task or, for example, many thousands of tasks. From step **1502(1)**, execution proceeds to step **1502(2)** of the customer-needs process **1502**.

At step **1502(2)**, it is determined whether the customer is a first customer in the set of active customers. If so, concurrently: (1) the reservation-accounting process **1506** is initiated; (2) the resource-allocation process **1508** is initiated; and (3) the customer-needs process proceeds to step **1502(3)**. If it is determined at step **1502(2)** that the customer is not the first customer in the set of active customers, execution proceeds to step **1502(3)**. At step **1502(3)**, it is determined whether a total number of active customers has changed. If not, the customer-needs process **1502** proceeds to step **1504(1)** of the reservations process **1504**. If it is determined at step **1502(3)** that the total number of active customers has changed, execution proceeds to step **1502(4)**.

25

At step **1502(4)**, an outer stationary distribution is computed for all active customers. The outer stationary distribution may be computed, for example, as described with respect to FIGS. 7 and 9. Accordingly, for each active customer, computing the outer stationary distribution involves computing, for each active customer, a customer-resource percentage according to a proportion of the current set of consumable resources that should be allocated to the active customer. All such customer-resource percentages should generally sum to one. In a typical embodiment, step **1502(4)** results in a balanced-utilization partitioning scheme as described with respect to FIGS. 6-10. More particularly, the balanced-utilization partitioning scheme includes one partition for each active customer. From step **1502(4)**, execution proceeds to step **1502(5)**.

At step **1502(5)**, an effective distribution is computed for each reservation. The effective distribution can be computed as described with respect to FIGS. 7 and 9. From step **1502(5)**, execution proceeds to step **1502(6)**. At step **1502(6)**, a flow-control clocking weight is calculated for each reservation. In a typical embodiment, the effective distribution is a primary component of each reservation's flow-control clocking weight. As described with respect to FIGS. 7, 9, 11, and 13, each reservation's flow-control clocking weight can be calculated as a product of the effective distribution and a total number of consumable resources in the current set of consumable resources. From step **1502(6)**, execution proceeds to step **1504(1)** of the reservations process **1504**.

The reservations process **1504** begins at step **1504(1)**. At step **1504(1)**, a new reservation is added to the priority queue. The new reservation corresponds to the customer-specific task group created at step **1502(1)** of the customer-needs process **1502**. From step **1504(1)**, execution proceeds to step **1504(2)**. At step **1504(2)**, an initial value is assigned to a wait-time for the new reservation. As described with respect to FIG. 13, the flow-control instance maintains wait-times for each reservation in the priority queue. The wait-times are considered a measure of how needy the reservations are for the class of consumable resources. The initial value of the wait-time for the new reservation can be, for example, a negative value of the new reservation's effective distribution (computed at step **1502(5)**). In various other embodiments, the initial value of the wait-time can be, for example, a negative value of the new reservation's flow-control clocking weight (computed at step **1502(6)**).

The reservation-accounting process **1506** begins at step **1506(1)**. At step **1506(2)**, the flow-control instance waits one clocking cycle (e.g., one second). From step **1506(2)**, execution proceeds to step **1506(3)**. At step **1506(3)**, the wait-time for each reservation in the priority queue is updated based on the reservation's flow-control clocking weight. The update yields a new value for the wait-time. In a typical embodiment, for each reservation, this update effects a clocking of the reservation in a negative direction as described with respect to FIG. 13. An example of how the wait-time for each reservation can be clocked negatively is described with respect to FIG. 14. From step **1506(3)**, execution proceeds to step **1506(4)**. At step **1506(4)**, the priority queue is reordered by wait-time, where a lowest wait-time indicates a most-needy reservation and a highest wait-time indicates a least-needy reservation. From step **1506(4)**, execution returns to step **1506(2)** and proceeds as described above.

The resource-allocation process **1508** begins at step **1508(1)**. At step **1508(2)**, it is determined whether any consumable resources in the current set of consumable resources are free (i.e., available for use). If not, the resource-allocation process **1508** remains at step **1508(2)** until a consumable resource in

26

the current set of consumable resources is free. If it is determined at step **1508(2)** that a consumable resource in the current set of consumable resources is free, execution proceeds to step **1508(3)**. At step **1508(3)**, the consumable resource is marked as "in use." From step **1508(3)**, execution proceeds to step **1506(5)** of the reservation-accounting process **1506**.

At step **1506(5)**, the consumable resource is assigned to a selected reservation in the priority queue that is qualified to use the consumable resource. In a typical embodiment, the selected reservation is the most-needy reservation as defined by having the lowest wait-time in the priority queue. At this point, the selected reservation is granted access to the assigned consumable resource until the selected reservation is finished, as described in greater detail below. From step **1506(5)**, execution proceeds to step **1504(6)** of the reservations process **1504**.

At step **1504(6)**, it is determined whether the selected reservation is still active. In a typical embodiment, the selected reservation is still active if the selected reservation includes at least one task that has not been processed to completion. If not, execution proceeds to step **1504(7)**. At step **1504(7)**, the selected reservation is removed from the priority queue. From step **1504(7)**, execution proceeds to step **1502(3)** and continues as described above. If it is determined at step **1504(6)** that the selected reservation is still active, execution proceeds to step **1504(3)**.

At step **1504(3)**, the selected reservation's wait-time is temporarily updated based on an average utilization time. The average utilization time can be represented in units of resource-time (i.e., as a product of resources and time). If wall time is utilized, the units of resource-time could, more specifically, be considered resource-seconds. The average utilization time is an average number of resource-time units (e.g., resource-seconds) that tasks grouped under the selected reservation's task key have utilized a consumable resource before yielding back the consumable resource. The average utilization time is based on previous consumable-resource assignments. If the selected reservation has not previously been assigned a consumable resource, the average utilization time may be a default value such as, for example, one resource-second. In a typical embodiment, a product of the average utilization time and the flow-control clocking weight is added to the selected reservation's wait-time. In a typical embodiment, this update prevents the selected reservation from remaining at the top of the priority queue. From step **1504(3)**, execution proceeds to step **1504(4)**.

At step **1504(4)**, it is determined whether a task of the selected reservation is finished with the assigned consumable resource. Execution remains at step **1504(4)** until a currently executing task is finished. As each task of the selected reservation is finished, execution proceeds to step **1504(5)**. At step **1504(5)**, the selected reservation's wait-time is adjusted based on an actual utilization time (i.e., resource-time) of the finished task. The adjustment typically involves subtracting the product temporarily applied at step **1504(3)** and, in its place, adding a product of the actual utilization time and one resource, thereby yielding units of resource-time (e.g., resource-seconds). For example, in various embodiments in which the actual utilization time is expressed in fractional units of wall time, the actual utilization time is precise to a nearest millisecond. In a typical embodiment, the average utilization time for the selected reservation is also updated based on the actual utilization time. From step **1504(5)**, execution proceeds to step **1504(6)** and proceeds as described above.

FIG. 16 illustrates a collection 1600 of interactive processes that may be executed by a flow-control instance such as, for example, the flow-control instance 200 of FIG. 2A. The flow-control instance manages a dynamically changing set of consumable resources relative to a managed task type requiring utilization of a class of consumable resources. In a typical embodiment, the collection 1600 is used to control sharing of the set of consumable resources among a dynamically changing set of active customers. The collection 1600 may be executed by a computer cluster such as, for example, the computer cluster 100 of FIG. 1.

In similar fashion to the collection 1500 of FIG. 15, the collection 1600 includes a customer-needs process 1602, a reservations process 1604, a reservation-accounting process 1606, and a resource-allocation process 1608. Differently than the collection 1500, the collection 1600 additionally includes a customer-priorities process 1603. In a typical embodiment, each process in the collection 1600 is executed in parallel by one or more computers of the computer cluster. Upon instantiation of the flow-control instance, no process of the collection 1600 is typically in execution. Execution of the collection 1600 typically begins with a first execution of the customer-needs process 1602. The customer-needs process 1600 begins at step 1602(1).

At step 1602(1), a customer-specific task group is created. In a typical embodiment, the customer-specific task group is created responsive to a request for service being received from a customer. Creation of the customer-specific task group typically involves grouping tasks, for example, by task key as described above with respect to FIGS. 3-5. It should be appreciated that the customer-specific task group might include a single task or, for example, many thousands of tasks. From step 1602(1), execution proceeds to step 1602(2).

At step 1602(2), it is determined whether the customer is a first customer in the set of active customers. If so, concurrently: (1) the reservation-accounting process 1606 is initiated; (2) the resource-allocation process 1608 is initiated; and (3) the customer-needs process proceeds to step 1602(3). If it is determined at step 1602(2) that the customer is not the first customer in the set of active customers, execution proceeds to step 1602(3). At step 1602(3), it is determined whether a total number of active customers has changed. If not, the customer-needs process 1602 proceeds to step 1604(1) of the reservations process 1604. If it is determined at step 1602(3) that the total number of active customers has changed, execution proceeds to step 1602(4).

At step 1602(4), an outer stationary distribution is computed for all active customers. The outer stationary distribution may be computed, for example, as described with respect to FIGS. 7, 9, and 11. Accordingly, for each active customer, computing the outer stationary distribution involves assigning, to each active customer, a customer-specific resource percentage according to a proportion of the current set of consumable resources that should be allocated to the active customer. All such customer-specific resource percentages should generally sum to one. From step 1602(4), execution proceeds to step 1603(1) of the customer-priorities process 1603.

The customer-priorities process 1603 begins at step 1603(1). At step 1603(1), an inner stationary distribution is computed for each active customer. Computation of the inner stationary distribution involves calculating a priority-based resource percentage for each priority grouping of the customer's active reservations. A priority grouping exists for each priority level at which the customer has active reservations. All such priority-based resource percentages should sum to

one. The inner stationary distribution may be computed, for example, as described with respect to FIGS. 11-12.

In a typical embodiment, step 1603(1) results in a balanced-utilization partitioning scheme as described with respect to FIGS. 6-12. More particularly, the balanced-utilization partitioning scheme includes one partition for each priority grouping so that a total number of partitions equals a number of priority groupings across all customers. From step 1603(1), execution proceeds to step 1603(2). At step 1603(2), an effective distribution is computed for each reservation. The effective distribution can be computed as described with respect to FIGS. 11-12. From step 1603(2), execution proceeds to step 1603(3).

At step 1603(3), a flow-control clocking weight is calculated for each reservation. In a typical embodiment, the effective distribution is a primary component of each reservation's flow-control clocking weight. As described with respect to FIGS. 7, 9, 11, and 13, each reservation's flow-control clocking weight can be calculated as a product of the effective distribution and a total number of consumable resources in the current set of consumable resources. From step 1603(3), execution proceeds to step 1604(1) of the reservations process 1604.

The reservations process 1604 begins at step 1604(1). At step 1604(1), a new reservation is added to the priority queue. The new reservation corresponds to the customer-specific task group created at step 1602(1) of the customer-needs process 1602. From step 1604(1), execution proceeds to step 1604(2). At step 1604(2), an initial value is assigned to a wait-time for the new reservation. As described with respect to FIG. 13, the flow-control instance maintains wait-times for each reservation in the priority queue. The wait-times are considered a measure of how needy the reservations are for the class of consumable resources. The initial value of the wait-time for the new reservation can be, for example, a negative value of the new reservation's effective distribution (computed at step 1603(2)). In various other embodiments, the initial value of the wait-time can be, for example, a negative value of the new reservation's flow-control clocking weight (computed at step 1603(3)).

The reservation-accounting process 1606 begins at step 1606(1). At step 1606(2), the flow-control instance waits one clocking cycle (e.g., one second). From step 1606(2), execution proceeds to step 1606(3). At step 1606(3), the wait-time for each reservation in the priority queue is updated based on the reservation's flow-control clocking weight. This update yields a new value for the wait-time. In a typical embodiment, for each reservation, this update also effects a clocking of the reservation in a negative direction as described with respect to FIG. 13. An example of how the wait-time for each reservation can be clocked negatively is described with respect to FIG. 14. From step 1606(3), execution proceeds to step 1606(4). At step 1606(4), the priority queue is reordered by wait-time, where a lowest wait-time indicates a most-needy reservation and a highest wait-time indicates a least-needy reservation. From step 1606(4), execution returns to step 1606(2) and proceeds as described above.

The resource-allocation process 1608 begins at step 1608(1). At step 1608(2), it is determined whether any consumable resources in the current set of consumable resources are free (i.e., available for use). If not, the resource-allocation process 1608 remains at step 1608(2) until a consumable resource in the current set of consumable resources is free. If it is determined at step 1608(2) that a consumable resource in the current set of consumable resources is free, execution proceeds to step 1608(3). At step 1608(3), the consumable

resource is marked as “in use.” From step **1608(3)**, execution proceeds to step **1606(5)** of the reservation-accounting process **1606**.

At step **1606(5)**, the consumable resource is assigned to a selected reservation in the priority queue that is qualified to use the consumable resource. In a typical embodiment, the selected reservation is the most-needy reservation as defined by having the lowest wait-time in the priority queue. At this point, the selected reservation is granted access to the assigned consumable resource until the selected reservation is finished, as described in greater detail below. From step **1606(5)**, execution proceeds to step **1604(6)** of the reservations process **1604**.

At step **1604(6)**, it is determined whether the selected reservation is still active. In a typical embodiment, the selected reservation is still active if the selected reservation includes at least one task that has not been processed to completion. If not, execution proceeds to step **1604(7)**. At step **1604(7)**, the selected reservation is removed from the priority queue. From step **1604(7)**, execution proceeds to step **1602(3)** and continues as described above. If it is determined at step **1604(6)** that the selected reservation is still active, execution proceeds to step **1604(3)**.

At step **1604(3)**, the selected reservation’s wait-time is temporarily updated based on an average utilization time. The average utilization time can be represented in units of resource-time (i.e., as a product of resources and time). If wall time is utilized, the units of resource-time could, more specifically, be considered resource-seconds. The average utilization time is an average number of resource-time units (e.g., resource-seconds) that tasks grouped under the selected reservation’s task key have utilized a consumable resource before yielding back the consumable resource. The average utilization time is based on previous consumable-resource assignments. If the selected reservation has not previously been assigned a consumable resource, the average utilization time may be a default value such as, for example, one resource-second. In a typical embodiment, a product of the average utilization time and the flow-control clocking weight is added to the selected reservation’s wait-time. In a typical embodiment, this update prevents the selected reservation from remaining at the top of the priority queue. From step **1604(3)**, execution proceeds to step **1604(4)**.

At step **1604(4)**, it is determined whether a task of the selected reservation is finished with the assigned consumable resource. Execution remains at step **1604(4)** until a currently executing task is finished. As each task of the selected reservation is finished, execution proceeds to step **1604(5)**. At step **1604(5)**, the selected reservation’s wait-time is adjusted based on an actual utilization time of the finished task. The adjustment typically involves subtracting the product temporarily applied at step **1604(3)** and, in its place, adding a product of the actual utilization time and ‘one’ resource, thereby yielding units of resource-time (e.g., resource-seconds). For example, in various embodiments in which the actual utilization time is expressed in fractional units of wall time, the actual utilization time is precise to a nearest millisecond. In a typical embodiment, the average utilization time for the selected reservation is also updated based on the actual utilization time. From step **1604(5)**, execution proceeds to step **1604(6)** and proceeds as described above.

For purposes of this disclosure, an information handling system may include any instrumentality or aggregate of instrumentalities operable to compute, calculate, determine, classify, process, transmit, receive, retrieve, originate, switch, store, display, communicate, manifest, detect, record, reproduce, handle, or utilize any form of information, intelligence,

or data for business, scientific, control, or other purposes. For example, an information handling system may be a personal computer (e.g., desktop or laptop), tablet computer, mobile device (e.g., personal digital assistant (PDA) or smart phone), server (e.g., blade server or rack server), a network storage device, or any other suitable device and may vary in size, shape, performance, functionality, and price. The information handling system may include random access memory (RAM), one or more processing resources such as a central processing unit (CPU) or hardware or software control logic, ROM, and/or other types of nonvolatile memory. Additional components of the information handling system may include one or more disk drives, one or more network ports for communicating with external devices as well as various input and output (I/O) devices, such as a keyboard, a mouse, touch-screen and/or a video display. The information handling system may also include one or more buses operable to transmit communications between the various hardware components.

Although various embodiments of the method and apparatus of the present invention have been illustrated in the accompanying Drawings and described in the foregoing Detailed Description, it will be understood that the invention is not limited to the embodiments disclosed, but is capable of numerous rearrangements, modifications and substitutions without departing from the spirit of the invention as set forth herein.

What is claimed is:

1. A method comprising:

empirically analyzing, by at least one process on a computer cluster comprising a plurality of computers, a set of active reservations and a current set of consumable resources belonging to a class of consumable resources; wherein the empirically analyzing comprises computing a flow-control clocking weight for each reservation of the set of active reservations;

wherein each active reservation is of a managed task type and comprises a group of one or more tasks requiring access to a consumable resource of the class;

based on the empirically analyzing, clocking the set of active reservations each clocking cycle;

responsive to the clocking, sorting, by the at least one process, a priority queue of the set of active reservations; wherein the sorting comprises sorting the priority queue by wait-times, the wait-times comprising a wait-time for each active reservation in the set of active reservations; and

wherein each wait-time comprises fractional units of resource-time.

2. The method of claim 1,

wherein each flow-control clocking weight comprises fractional units of resources.

3. The method of claim 2, wherein, for each reservation of the set of active reservations, the wait-time is given an initial value upon creation of the reservation, the initial value comprising a factor of the flow-control clocking weight.

4. The method of claim 3, wherein the clocking comprises, each said clocking cycle, decrementing the wait-time for each reservation in the set of active reservations by a factor of the flow-control clocking weight.

5. The method of claim 4, wherein the factor is a product of the flow-control clocking weight, an actual average capacity of the current set of consumable resources during the clocking cycle, and a length of the clocking cycle.

6. The method of claim 2, comprising, responsive to a determination that at least one consumable resource in a current set of consumable resources is free, assigning the at

31

least one consumable resource to a reservation of the set of active reservations that is at a top of the priority queue.

7. The method of claim 6, comprising, as each task of the reservation finishes accessing the assigned consumable resource, incrementing the wait-time by a factor of resource-time units utilized by the task.

8. The method of claim 2, wherein the wait-times are maintained in a double-precision floating-point format.

9. The method of claim 2, comprising:
wherein the clocking cycle is defined in units of wall time;
and
wherein the wait-times comprise fractional units of resource-seconds.

10. The method of claim 1, wherein the empirically analyzing is repeated as new reservations of the managed task type are created on the computer cluster.

11. The method of claim 1, wherein the class of consumable resources is access to a bottleneck point of a business process.

12. The method of claim 1, wherein the class of consumable resources comprises at least one selected from the group consisting of: floating software licenses and voice ports.

13. The method of claim 1, wherein the method is performed relative to multiple managed task types.

14. The method of claim 1, wherein at least a portion of the current set of consumable resources are external to the computer cluster.

15. An information handling system, comprising:
a computer cluster comprising a plurality of computers,
wherein the computer cluster is operable to:
empirically analyze a set of active reservations and a current set of consumable resources belonging to a class of consumable resources;
wherein the empirical analysis comprises computing a flow-control clocking weight for each reservation of the set of active reservations;
wherein each active reservation is of a managed task type and comprises a group of one or more tasks requiring access to a consumable resource of the class;
based on the empirical analysis, clock the set of active reservations each clocking cycle;
responsive to the clocking, sort a priority queue of the set of active reservations;
wherein the sort comprises sorting the priority queue by wait-times, the wait-times comprising a wait-time for each active reservation in the set of active reservations; and
wherein each wait-time comprises fractional units of resource-time.

32

16. The information handling system of claim 15, wherein each flow-control clocking weight comprises fractional units of resources.

17. The information handling system of claim 16, comprising:
wherein, for each reservation of the set of active reservations, the wait-time is given an initial value upon creation of the reservation, the initial value comprising a factor of the flow-control clocking weight,
wherein the clocking comprises, each said clocking cycle, decrementing the wait-time for each reservation in the set of active reservations by a factor of the flow-control clocking weight.

18. The information handling system of claim 17, wherein the computer cluster is operable to:
responsive to a determination that at least one consumable resource in a current set of consumable resources is free, assign the at least one consumable resource to a reservation of the set of active reservations that is at a top of the priority queue.

19. The information handling system of claim 17, wherein the computer cluster is operable to:
as each task of the reservation finishes accessing an assigned consumable resource, increment the wait-time by a factor of resource-time units utilized by the task.

20. A computer-program product comprising a non-transitory computer-usable medium having computer-readable program code embodied therein, the computer-readable program code adapted to be executed to implement a method comprising:

empirically analyzing a set of active reservations and a current set of consumable resources belonging to a class of consumable resources;
wherein the empirically analyzing comprises computing a flow-control clocking weight for each reservation of the set of active reservations;
wherein each active reservation is of a managed task type and comprises a group of one or more tasks requiring access to a consumable resource of the class;
based on the empirically analyzing, clocking the set of active reservations each clocking cycle;
responsive to the clocking, sorting a priority queue of the set of active reservations;
wherein the sorting comprises sorting the priority queue by wait-times, the wait-times comprising a wait-time for each active reservation in the set of active reservations;
and
wherein each wait-time comprises fractional units of resource-time.

* * * * *